

VARIABLE LENGTH PIPELINE WITH PARALLEL FUNCTIONAL UNITS

Inventors: Paul W. DeMone
Peter B. Gillingham

Exp. Mail No.: EL13620306045

Attorney Docket No.: SLDM1029WSW
/wsw/sldm/1029.app

VARIABLE LENGTH PIPELINE WITH PARALLEL FUNCTIONAL UNITS

Inventor: Paul W. DeMone
Peter B. Gillingham

5 BACKGROUND TO THE INVENTION

Field Of The Invention

The present invention relates generally to command processing applications in high bandwidth memory systems.

10 Cross Reference to Other Applications

The following pending application is owned by the assignee of the present application, and its contents are hereby incorporated by reference:

Serial No. 09/132,158 [Attorney Docket No. SLDM1025]
15 filed August 10, 1998, invented by Gustavson et. al and entitled, *MEMORY SYSTEM HAVING SYNCHRONOUS-LINK DRAM (SLDRAM) DEVICES AND CONTROLLER*

Description of the Related Art

20 The evolution of the dynamic random access memories used in computer systems has been driven by ever-increasing speed requirements mainly dictated by the microprocessor industry. Dynamic random access memories (DRAMs) have generally been the predominant memories used for computers due to their optimized storage capabilities. This large
25 storage capability comes with the price of slower access time and the requirement for more complicated interaction

between memories and microprocessors/microcontrollers than in the case of say static random access memories (SRAMs) or non-volatile memories.

In an attempt to address this speed deficiency, DRAM design has implemented various major improvements, all of which are well documented. Most recently, the transition from Fast Page Mode (FPM) DRAM to Extended Data Out (EDO) DRAMs and synchronous DRAMs (SDRAMs) has been predominant. Further speed increases have been achieved with double data rate (DDR) SDRAM, which synchronizes data transfers on both clock edges. New protocol based memory interfaces have recently been developed to further increase the bandwidth and operating frequencies of synchronous memories.

As the complexity of these memories has increased, the associated control systems responsible for internally managing the operation of the memories have also become more complex. These command-driven control systems internally must typically process a stream of commands or instructions that overlap in execution time and have programmable latency (time from receipt of command to first control outputs asserted in response). Programmable latency is desirable in such systems in order to allow the memory controller to schedule the use of shared data, address or control buses for optimum usage. Since the processing of two or more commands may be required to occur

simultaneously, many control systems implement multiple functional units operating in parallel. The minimum latency of the control system is therefore limited by the need to (i) decode the command control field(s), (ii) 5 determine the programmed latency associated with the identified command, and (iii) issue the command to a number of parallel functional units before the first control output action can be determined for use by the memory.

A conventional implementation of such a memory 10 system control block comprises a single front end decoding block which decodes external commands and issues internal commands to multiple identical functional elements capable of operating in parallel. The minimum latency therefore cannot be shorter than the time it takes to decode the 15 command in the front end block plus the time required to issue the command to a parallel functional unit, and finally, the time that the functional unit takes to initialize and issue its first control action. The common approach to reducing the minimum latency described above is 20 by replicating the command decoding logic within each parallel functional unit and feeding the command stream to all parallel functional units simultaneously to eliminate the issue and initialization delay. This advantage comes with the cost of a large increase in overall logic 25 complexity, redundant logic, and increased power

consumption. As frequency and bandwidth requirements increase, there is a need for a memory system control block which makes optimum use of area and power consumption and which can process commands with a reduced minimum latency than previously achieved in the prior art.

SUMMARY OF THE INVENTION

It is therefore an object of the present invention to provide a command processing system for use in a high bandwidth memory interface which processes commands with a minimum latency.

It is another object of the present invention to provide the command processing system with a minimum increase to the command circuitry.

According to the invention, roughly described, a packet-driven memory control system which implements a variable length pipeline includes a command front end and one or more parallel command sequencers. The command front end decodes an external command packet into an internal command and issues it to a selected one of the command sequencers. The command has associated therewith a desired latency value. A first group of one or more memory control steps for the given command is performed by the command front end if the desired latency value is less than a threshold latency value, or by the selected command

sequencer if the desired latency value is greater than or equal to the threshold latency value. The remainder of the memory control steps required for the command are performed by the selected command sequencer. If the first control
5 steps are to be performed by the selected command sequencer, then depending on the desired latency value, the command sequencer further may insert one or more wait states before doing so.

10 **BRIEF DESCRIPTION OF THE DRAWINGS**

Fig. 1 is a simplified block diagram of a synchronous link memory system incorporating features of the invention.

15 Figs. 2A and 2B together represent a simplified block diagram of the synchronous link DRAM (SLDRAM) module or integrated circuit of Fig. 1.

Figs. 3A and 3B are schematics and timing diagrams illustrating the conversion of external command, address and flag signals of Fig. 2 into internal command and
20 address signals to be processed.

Fig. 4 is a simplified block diagram of the command processing pipeline incorporating an embodiment of the invention.

25 Fig. 5A is a conceptual diagram illustrating the processing of a minimum latency page read command.

Fig. 5B is a conceptual diagram illustrating the processing of a non-minimum latency page read command.

Fig. 6 is a block diagram of a command sequencer according to an embodiment of the invention.

5 Fig. 7 is a block diagram of a control signal variable delay circuit with a delay resolution shorter than one clock period.

10 Fig. 8 is a block diagram illustrating input and output circuits connected to shared bus signal lines CTL/ADDR in Fig. 4.

Fig. 9 is a timing diagram illustrating the operation of the circuits of Fig. 8.

DETAILED DESCRIPTION OF THE INVENTION

15 Fig. 1 provides a simplified view of a memory system employing a packet based synchronous link architecture (SLDRAM). The system which is described more fully in the above-incorporated patent application generally comprises a command module 150 (typically implemented by a memory
20 controller), and a plurality of SLDRAM modules or individual IC's 110-180. The command link 151 is used by the command module to issue commands, a command system clock, control signals and address information to each of the SLDRAMs. Data is written into or read out of the
25 SLDRAMs via the DataLinks 155a and 156a in synchronization

with source-synchronous clocks 155b, 155c, 156b and 156c. Within this system, an embodiment of the command processing in accordance with the invention will be described.

Fig. 2A and 2B together illustrate the general structure of an SDRAM memory integrated circuit of Fig. 1. The structure and operation of the circuit is described broadly in the above-incorporated patent application. The command decode and sequencer unit 504 will be described in more detail below.

Fig. 3A illustrates the input stage of the command decoder 504 of Fig. 2A. The incoming external command and address signals CA[9:0] along with the FLAG and command clock CCLK signals are received via input cells each comprising an input protection device 50, an input buffer 51 and two D-type flip/flops 52 and 53 for latching the command/address and FLAG signals on both rising and falling edges of the command clock CCLK. As a result, the eleven (11) incoming signals made up of FLAG and CA[9:0] operating at 400Mbps are converted internally into twenty two (22) internal command/address signals consisting of FLAG_R, FLAG_F, CA_R[9:0], and CA_F[9:0], operating at 200Mbps. The command clock also has a delay locked loop (DLL) and vernier delay in its input path, which are used to properly latch the incoming commands and address signals at the appropriate time within the system.

Fig. 3B illustrates the relative timing of the input stage. CCLK is a free-running clock. Upon assertion of the FLAG signal, command/address words begin to be latched on the rising edge of the delayed internal version of the command clock CCLKH. On the subsequent rising edge of the internal flag signal FLAG_R, the internal command/address words begin to be accepted into the system at one half the frequency of the external CCLK. The command/address words are alternated between the rising and falling edge command/address internal busses CA_R[9:0] and CA_F[9:0] as indicated by A0, A1, A2, A3, etc.

Fig. 4 is a block diagram illustrating the command processing pipeline according to an embodiment of the invention. A command decoder front end CFE 200 receives the command packet as four consecutive 10-bit words on CA[9:0]. It then internally assembles and decodes the command packet into a 31-bit internal command COM[30:0] which is issued to a selected one of a plurality of parallel functional units or command sequencers 201-208. The CFE 200 also generates a 6-bit command delay signal COMDEL[5:0] which is determined by comparing the latency in the selected latency register with a predetermined threshold. The CFE 200 initializes each of the sequencers by asserting the ISSUE0-ISSUE7 signals. The available or busy state of a sequencer is feedback to the CFE via the

BUSY0-BUSY7 signals. Both the CFE 200 and each of the sequencers also have a multi-bit control/address output CTL/ADD which is used to send out the control signals to the memory banks, the data path, etc. The CTL/ADD signal coming from the CFE 200 corresponds to control signals being generated by the CFE itself as will be described in more detail below.

With reference to Fig. 2A and Fig. 4, in accordance with an embodiment of the present invention, an SLDRAM memory device receives streams of command packets as 4 consecutive 10-bit words on the CA[9:0] bus. Each 40 bit command packet is assembled and then decoded by the command front end or CFE block 200. For SLDRAM commands that utilize user-programmable latencies (such as memory reads, memory writes, and register reads), the CFE 200 selects the appropriate latency value based upon the command type and issues the command packet and latency value to one of eight identical parallel functional units called command sequencers 201-208, with sequencer 0, 201 having the highest priority and sequencer 7, 208 having the lowest priority. The determination of whether to perform the first group of control steps within the CFE 200 or to forward the entire command to a selected command sequencer depends on the command's specified latency. Once a command is decoded, if the desired latency is determined to be

shorter than a predetermined threshold, then the CFE 200 executes the first several control steps using control logic located within the CFE block 200, and simultaneously issues the command to a parallel functional unit and initializes that unit. Subsequently, the control action sequence is seamlessly taken over by the selected command sequencer which recognizes (based on the latency value accompanying the command) that the CFE 200 has already performed the initial control actions. The selected command sequencer therefore skips performing these actions which it would normally do, and instead proceeds directly to execute the remaining control actions necessary to process the command. For example, if a page read command is dispatched by the system's command module 150, the CFE 200 within a particular SLD RAM recognizes this command as a special case if the page read latency register located in the SLD RAM device was programmed to the minimum value by a previous command. When this occurs, special logic in the CFE 200 performs the first two control actions for a page read (column open (select) for the low array, and initiation of a read operation within the data output path) simultaneously with issuing the command to an idle sequencer. The sequencer itself is designed to recognize the special case of a minimum latency page read and will skip the first two control steps performed by the special

logic in the CFE 200 and instead directly proceed to the remaining instructions to complete the page read command.

Fig. 5A illustrates the relative timing of processing of a minimum latency page read operation by the CFE 200 and a selected command sequencer. Upon a predetermined rising edge of the command clock, CCLKH, the CFE opens a selected column, initiates the data path (precharging and equalizing of data buses) and issues the page read command to the available sequencer, all within the first CCLK period. Then, during a subsequent CCLK period, the command sequencer opens a second column, column high (previously column low was opened) to initiate the read of the second column. Prior to the end of this subsequent clock cycle, the read data path begins to receive the low data bits corresponding to the low column which was opened by the CFE in the first CCLK cycle. Subsequently, after some delay, the read data path receives the high data bits corresponding to the high column which was opened by the sequencer, as described above. In this fashion, the labour (reading column low and column high) was divided between the CFE and a selected sequencer by having the CFE perform the first portion of the operation, and the sequencer perform the remaining portion of the operation.

If on the other hand, the desired latency is determined to be greater than a predetermined threshold i.e. if the actual page read latency register is programmed to a value greater than the minimum latency, then the CFE
5 200 executes none of the control actions and instead forwards them all to an available command sequencer. In this case, the selected sequencer also recognizes that the page read latency is greater than minimum and performs all control actions to accomplish the page read (after
10 inserting any necessary wait states).

Fig. 5B illustrates the relative timing of processing of a non-minimum latency page read operation by the CFE 200 and a selected command sequencer. Upon a predetermined rising edge of the command clock, CCLKH, the
15 CFE 200 recognizes that the requested command is a non-minimum latency command by the value written into the latency register, and immediately issues the command to an available command sequencer within the first CCLK period. The selected sequencer is initialized and a number of
20 latency states are inserted depending on the value in the page read delay register. Once the latency wait states have elapsed, the sequencer proceeds to execute the command in a manner similar to that described in Fig. 5A, i.e. a column open low is performed along with the initialization
25 of the data path. Subsequently, a column open high read

operation is performed during the second clock cycle and during that same cycle, low data starts to appear on the read data path. Finally, an optional row close command is executed during the third clock cycle and the high data appears on the read data path. By optionally performing the initial page read control actions simultaneously with the issuing of the page read command to an idle sequencer, the minimum page read latency is reduced by one clock cycle.

In general, the CFE block 200 has the following procedure for receiving and processing a command:

-
- Assemble a 40-bit command packet when FLAG is asserted
 - Compare packet ID with the device ID in order to determine whether packet is heading to correct device
 - Decode 6-bit command field to:
 - Determine the command type (buffered, immediate, etc.)
 - Determine command latency
 - Issue command if all the following conditions are satisfied:
 - Command field contains a valid opcode
 - Command ID matches device ID

- FLAG protocol is obeyed, i.e. FLAG bit is asserted for one clock tick only (i.e. half a period)
- Command processing mode is enabled
- An idle command sequencer is available

Fig. 6 illustrates a command sequencer 321 corresponding with one of the command sequencers 201-208 in Fig. 4, according to an embodiment of the invention. The command signals COM[30:0] are received by a latch 300 which is enabled by a signal G from a Idle/Active module select block 303 (for a more detailed breakdown of the command packet, see Table 2.0 in the above-incorporated patent application). The output of the latch 300 is broken down into bank address signals BNK[2:0], register address signals REGA[3:0], column address signals COLA[6:0] and the actual command instructions CMD[5:0]. The BNK[2:0] signals are decoded by a 3-to-8 decoder 304 and then fed into output buffers 314 for high and low column block addresses signals YBKLO[7:0] and YBKHI[7:0], as well as being input into a miscellaneous decoder 317 for closing an open row RCLOSE[7:0]. The register addresses REGA[3:0] are output via buffers 315, while the column addresses COLA[6:0] are latched and then output via buffers 316; the LSB COLA[0] is optionally inverted by an LSB inverter 305 for performing

the second half of the word burst operation. The misc.
decoder 317 also receives the command instruction signals
CMD[5:0] as inputs. The required command latency delay is
input into the sequencer via lines COMDEL[5:1] into a 5-bit
5 counter 301 and with the least significant bit COMDELO
input into a latch 302. The counter and latch 301 and 302
respectively, also receive the G control signal from the
Idle/Active module select block 303. The output of the
counter 301 feeds into read latency decoders 360, read
10 command decoders 307, write latency decoders 308 and write
command decoder 309. If the sequencer is available, the
Idle/Active module select block 303 generates and asserts
an ACTIVE signal in response to an asserted ISSUE signal
from the CFE. The ACTIVE signal in turn enables the decoder
15 combining circuitry, AND gates 310 and 311. OR gate 312
selects between read and write command decoder outputs from
310 or 311 respectively to initiate a column operation via
block 319. The column operation block 319 also produces a
control signal which is used to control the buffers 314,
20 315 and 316, and also produces the output the control
signals COLENLO, COLENHI for internally enabling the
selected columns within the device. If a read command is
decoded along with its corresponding latency via 306, 307
and 310, a data output path command encoder 318 is used to
25 generate the data path output control signals DPO[4:0]. If

a write command is decoded along with its corresponding latency via 308, 309 and 311, a data input path command encoder 320 is used to generate the data path input control signals DPI[4:0]. The data path output and input command
5 decoders 318 and 320 are also controlled by the LSB from latch 302.

The sequencer 321 is one of eight identical functional units 201-208 as illustrated in Fig. 4. There is no interlocking between the sequencers or between a particular sequencer and the CFE 200. Therefore, the
10 command module (memory controller) must be aware of the actual delay values and schedule commands appropriately. The sequencer performs any one the following operations:

-
- 15 • all bank read/write commands except for row open which is performed by the command front end CFE
- all page read commands unless actual delay is programmed to minimum, in which case the CFE performs the data path initiate and part of the
20 column open
- all register read and read synch commands unless page read actual delay is programmed to minimum, in which case the CFE performs the data path initiate
-

As a further clarification as to the Division of Labour between the CFE and the command sequencers, Tables 1A, 1B and 1C are included below. These tables set forth the memory control steps performed by the CFE or by a sequencer, as the case may be, in response to a received command. As used herein, a "memory control step" is a step which drives the operation of a DRAM bank in a desired manner. The memory control steps set forth in Tables 1A, 1B and 1C are illustrative ones of such steps which are used in the present embodiment.

Table 1A: Division of Labor - Read Operations

Command	Command Front End Memory Control Steps	Sequencer Memory Control Steps
Read Page (BURST4)**	If latency = minimum <ul style="list-style-type: none"> open column low, initiate DPO transfer issue command to sequencer	If latency > minimum <ul style="list-style-type: none"> insert necessary wait states open column low, initiate DPO transfer open column high optional precharge
Read Page (BURST8)***	If latency = minimum <ul style="list-style-type: none"> open column low, initiate DPO transfer issue command to sequencer	If latency > minimum <ul style="list-style-type: none"> insert necessary wait states open column low, initiate DPO transfer open column high open column low*, initiate DPO transfer open column high* optional precharge
Read Bank (BURST4)	open row issue command to sequencer	insert necessary wait states open column low, initiate DPO transfer open column high optional precharge
Read Bank (BURST8)	open row issue command to sequencer	insert necessary wait states open column low, initiate DPO transfer open column high open column low*, initiate DPO transfer open column high* optional precharge

* LSB of column address is complemented.

** BURST4 refers to a burst of 4 consecutive 18-bit data words.

*** BURST8 refers to a burst of 8 consecutive 18-bit data words.

Table 1B: Division of Labor - Write Operations

Command	Command Front End Memory Control Steps	Sequencer Memory Control Steps
Write Page (BURST4)	<ul style="list-style-type: none"> • issue command to sequencer 	<ul style="list-style-type: none"> • insert necessary wait states • initiate DPI transfer • open column low • open column high • optional precharge
Write Page (BURST8)	<ul style="list-style-type: none"> • issue command to sequencer 	<ul style="list-style-type: none"> • insert necessary wait states • initiate DPI transfer • open column low • open column high, initiate DPI transfer • open column low* • open column high* • optional precharge
Write Bank (BURST4)	<ul style="list-style-type: none"> • open row • issue command to sequencer 	<ul style="list-style-type: none"> • insert necessary wait states • initiate DPI transfer • open column low • open column high • optional precharge
Write Bank (BURST8)	<ul style="list-style-type: none"> • open row • issue command to sequencer 	<ul style="list-style-type: none"> • insert necessary wait states • initiate DPI transfer • open column low • open column high , initiate DPI transfer • open column low* • open column high* • optional precharge

*LSB of column address is complemented

Table 1C: Event Operations

Command	Command Front End Memory Control Steps	Sequencer Memory Control Steps
Read Register	If latency = minimum • initiate DPO transfer (register) issue command to sequencer	If latency > minimum • insert necessary wait states • initiate DPO transfer (register) drive address to register selection MUX
Read Sync	If latency = minimum • initiate DPO read sync issue command to sequencer	If latency > minimum • insert necessary wait states • initiate DPO read sync
Row Open	open row	
Row Close	close row	
Register Write, Event, Stop Read Sync, Drive DCLKs, Disable DCLKs	issue command to immediate command block	

In general, the command sequencers perform bank read and writes, page writes and all the rest of operations with a programmable latency which is not set to a minimum value. It will be appreciated that the CFE and the sequencers never perform the same control step at the same time, the memory controller being responsible for scheduling instructions in such a way that the CFE and sequencers will not be generating control signals which create contention on the CTL/ADDR bus of Fig. 4. Similarly, the memory controller is responsible for ensuring that the parallel sequencers do not create contention. Note that two parallel sequencers can operate simultaneously and still

not create contention if, for example, they are generating control signals for different banks of memory controlled by different signal lines.

The command pipeline described above gives rise to one timing outcome which must be compensated. Namely, since command latencies are programmed in increments of clock "ticks" i.e. half clock cycles, and the command pipeline operates with a full clock period (i.e. 2 ticks), for latencies requiring an odd number of delays, a mismatch arises between the latency ticks and the command clock period, since the command pipeline cannot insert the appropriate number of tick delays based solely on its clock period. For an even number of delays, there is no mismatch between the number of delays required and the command pipeline clock period. As a result, a method for inserting an additional tick delay for odd-numbered latencies is implemented in a preferred embodiment of the invention, as will be discussed below.

More generally, in order to generate control signals with timing resolution T_{res} using conventional synchronous logic design techniques, it is necessary to clock the logic with a clock period shorter than or equal to T_{res} . For a high timing resolution system (i.e. short T_{res}), this requires a high operating frequency for the control logic, resulting in relatively high power consumption, especially

in CMOS implementations due to the CV^2f term, and also resulting in the minimum timing resolution T_{res} being limited by the maximum operation frequency of the synchronous control logic. Conventional approaches to resolving this issue included simply designing the control logic to operate at the frequency necessary for the desired timing resolution T_{res} by use of an increased control pipeline depth or the use of special circuit level design techniques such as dynamic storage elements to achieve the desired frequency/resolution. However, as the operating frequencies have increased, simply forcing the control logic by design to operate at those frequencies is becoming more and more challenging.

According to a preferred embodiment of the invention, a half-period adjust scheme is implemented to address this timing resolution drawback. The control logic is designed to operate with a clock period that is an integral multiple N of the desired timing resolution T_{res} , i.e., the control logic operates with clock period $T_{cp}=N \times T_{res}$. As a result, control signal timing is represented in terms of an integral number P of T_{cp} clock periods plus a fraction F/N where F is an integer between 0 and $N-1$, $t_{cs} = (P + F/N) \times T_{cp}$. The implementation of the control logic to handle this timing is as follows:

1) Store the parameter F while using P to count out the desired number of clock periods.

2) Upon completion of P synchronous counting steps, use the parameter F to generate the output signal delayed from the logic clock by $(F/N) \times T_{cp}$.

One possible implementation is to use the parameter F to control the insertion of appropriately scaled delay elements within the signal path of the output control signal. An alternate implementation is to pass the parameter F to the functional logic being controlled for the delay to be effected there.

Specifically, with respect to the command processing described earlier, with command latencies programmed in ticks, and the command pipeline operating with a full clock period, (in this case 5ns.) the half-clock adjust solution according to the preferred embodiment of the invention consists of implementing the latency within the command pipeline to within the nearest clock count, or effectively dividing the tick count by two, and then adjusting for the final fraction portion according to the number of tick delays required by the latency. In the case of an even tick count latency, the resulting tick count implemented in the command pipeline is equivalent to the tick count programmed. For an odd tick count latency delay, the command pipeline delay ends up being early by a half a

clock period. In order to compensate for this effect, the command is flagged as requiring a "half-period adjustment" and the data path introduces an extra half clock delay.

Fig. 7 illustrates a general implementation of this aspect of the invention. A latency value is input along with a command and stored in a latency register, in this case a 6-bit unsigned value. For a read operation for example, the read latency associated with the read operation is processed as follows:

1) the control logic takes the upper 5 bits of the latency value and inserts that number of 5ns wait states within the command pipeline;

2) the least significant bit of the programmed latency value is passed along through the command pipeline as the "half clock adjust bit". When the wait states inserted in the command pipeline are completed the control logic asserts a control signal to the data output path logic along with the half clock adjust bit. If the half clock adjust bit is logic 1, then the data path further delays the read data by 2.5ns, alternately, if the half clock adjust bit is logic 0, then the data path does not insert any additional delay.

In general, the half period adjust scheme can be extended as follows. For a system with desired timing resolution T_{res} , the control pipeline can be clocked with a

clock with a period T_{cp} that is an integral multiple N that is a power of two times T_{res} , i.e., $T_{cp}=N \times T_{res}$, $N=2^n$. Referring to Fig. 7, a timing parameter is then represented as a binary M -bit fixed point value with the least significant n bits as a fraction of the T_{cp} clock period. The m timing parameter bits above the least significant n bits specify the synchronous logic delay count P . These bits are loaded into a down counter 710. The least significant n bits carry the fractional delay value F , and are loaded into a latch 712 for temporary storage. After the down counter 710 counts down P clock pulses, a zero detector 714 asserts the desired control signal. This control signal is provided to $N-1$ delay elements 716.1, 716.2, ..., 716. $N-1$ (collectively 716), which delay the control signal by respective amounts $1/N T_{cp}$, $2/N T_{cp}$, ..., and $(N-1)/N T_{cp}$. The control signal is also provided to one input of a multiplexer 718, as are the outputs of each of the delay elements. The n low order bits of the delay value are provided from the latch 712 to the select input of multiplexer 718. Thus the control signal, already delayed by P clock periods T_{cp} by the counter 710, is then further delayed by the specified fractional part F/N of a clock period by the delay elements 716 and multiplexer 718.

In the embodiment described herein, $M=6$, $m=5$, $n=1$, and $N=2$. In this case the control pipeline is clocked with

the clock period T_{cp} which is $N=2$ times the desired timing resolution T_{res} . The least significant $n=1$ bit of the timing value is therefore used to control a 2-to-1 multiplexer 718 to select the synchronous pipeline output signal delayed by 0 or $\frac{1}{2} T_{cp}$ as the control signal output by the control logic. In another example, with $n = 3$, the control pipeline is clocked with the clock period T_{cp} which is $N=8$ times the desired timing resolution T_{res} . The least significant 3 bits of a timing value are therefore used to control an 8-to-1 multiplexer 718 to select the synchronous pipeline output signal delayed by 0, $\frac{1}{8} T_{cp}$, $\frac{2}{8} T_{cp}$, $\frac{3}{8} T_{cp}$, $\frac{4}{8} T_{cp}$, $\frac{5}{8} T_{cp}$, $\frac{6}{8} T_{cp}$, or $\frac{7}{8} T_{cp}$ as the control signal output by the control logic.

Note that other implementations are possible within the scope of this aspect of the invention. For example, the delay elements 716 could be replaced if desired by a single delay line having $N-1$ taps. As another example, the delay elements 716 and the multiplexer 718 in combination could be replaced by a single variable delay element. Other variations will be apparent.

Contention-Free Signaling Scheme

As shown in Fig. 4, many of the CTL/ADDR leads that are driven by the CFE 200 or any of the command sequencers 201-208 are shared. Thus at different times they might be

driven by different controlling units. As mentioned above, the memory controller is responsible for ensuring, through proper scheduling of memory control steps, that no two of the controlling units assert signals on the same control line in the same clock pulse. Though not required in different embodiments, the memory module of the present embodiment uses a transition-based contention-free signaling scheme in order to achieve enhanced contention-free operation.

Fig. 8 is a block diagram illustrating the circuits connected immediately to one of the shared control lines 210-X. Control units 830 and 831 represent two units from the group consisting of the CFE 200 and the command sequencers 201-208, any of which can drive the control line 210-X. Functional unit 832 represents any of the functional units which receive commands from the shared bus, such as DRAM banks and data paths in and out. The bus holder cell 833 could be physically part of any of the control units or functional units, or could be physically a separate cell as shown in Fig. 8. The function of the bus holder cell 833 is described below.

Fig. 8 shows the output driver portion of the control units 830 and 831. Referring to control unit 830, the output driver comprises two D-type flip/flops 836 and 834 as well as a tri-state buffer 835. Flip-flop 836

receives a command "assert-X" at its D-input and the system clock CLK at its clock input and outputs, on its Q output, a control signal to enable the tri-state buffer 835. Flip/flop 834 receives the output of the tri-state buffer 835 at its D-input and CLK on its clock input and outputs its Q\ ("Q-not") output to the input of the tri-state buffer 835. The resulting output signal from control unit 830 is therefore the output of the tri-state buffer 835. A similar output driver structure exists for control unit 831 as illustrated in Fig. 8.

The bus holder cell 833 consists of two cross-coupled inverters 843 and 844 which essentially act as a shared SRAM (static random access memory) bit storing the most recently asserted value on control signal line 210-X, until overwritten. The output of each inverter is connected to the input of the other, and the output of inverter 843 (input of inverter 844) is connected to the shared signal line 210-X. The inverter 843 is designed with weak driving characteristics so it can be easily overcome with an opposite polarity signal driven onto the shared control line 210-X by one of the control units 830 or 831.

The input portion of functional unit 832 comprises two D-type flip/flops 840 and 841 and an exclusive OR (XOR) gate 842. Shared control signal 210-X is input into one of the inputs of the XOR gate 842 as well as to the D-input of

flip/flop 840, which in turn is clocked by the system clock CLK. The Q output of flip/flop 840 is input as the second input to the XOR gate 842, which then outputs to the D-input of flip/flop 841. The Q output of flip/flop 841 represents an "asserted-X" control signal within the functional unit which is used to implement some control operation in the functional unit 832.

Fig. 9 is a timing diagram illustrating the operation of the circuits of Fig. 8. Referring to Fig. 9, prior to an arbitrary system clock cycle, cycle 1 for example, initiated at sampling time t_0 , control unit 830 evaluates a command and decides to assert the corresponding control signal onto the shared control signal line 210-X. Since this is a fully synchronous system, control unit 830 will assert its request and upon the next rising edge of CLK at time t_1 , and after a short time delay, the control signal 210-X will experience a transition in its logic state from a logic low to a logic high (note that prior to this change, at sampling time t_1 , the shared control signal 210-X had a logic low value). The state of control signal 210-X is maintained by the bus holder cell 833 for the duration of cycle 1 until it is overwritten in the next cycle. At the end of cycle 1 and the beginning of cycle 2, control unit 831 evaluates a command action and chooses to assert X. At the end of cycle 2, demarcated by sampling

time t_2 , the shared control signal 210-X is still logic high, and therefore a state transition is detected by the functional unit 832. The D-flip/flop 840 stores the last 210-X value (output of 840) and the XOR gate 842 compares the current value of 210-X and last 210-X. Since at t_2 , 210-X is logic high and last 210-X is logic low, the X-asserted output of D-flip/flop 841 is made logic high by the XOR gate 842 on the rising clock edge. The functional unit 832 then proceeds to execute the control steps associated with the X-asserted control signal (not shown). Subsequently, during a third clock cycle, cycle 3, control unit 831 decides to continue to assert X. At time t_3 , the functional unit 832 samples the shared signal 210-X and finds it to be logic low, thereby indicating another state transition since sampling time t_2 . Since the command to continue to assert X was provided during cycle 3, 210-X will again change states after sampling time t_3 , and the last 210-X in the functional unit D-flip/flop 840 will also change states. However, since both 210-X and last 210-X still remain opposite in phase, the X-asserted output remains logic high, through the XOR action of 842. As can be seen from Fig. 9, three clock cycles are required between the time when an action is evaluated by a controlling unit and the time when an asserted control signal results in the functional unit. Also, from Fig. 9,

it can be seen that the two controlling units used to illustrate the operation, units 830 and 831 did not have to contend for the control signal 210-X bus over consecutive clock cycles. The system can continue to operate in this fashion with alternating control between control units on every clock cycle.

Alternate Embodiments and Applications

As higher clock frequencies will be required in future applications, deeper pipelining will also be required, and according to an embodiment of the invention, two or more clock cycles of control activity are selectively moved up into the command front end based on early or partial decoding of certain commands. For example, consider a case where commands require 4 consecutive pipeline stages D1, D2, D3 and D4 to completely capture, decode and issue to a parallel function unit (sequencer), illustrated in Table 2. The commands themselves take control actions C1, C2, and C3 in three consecutive clock cycles to perform. Without the division of the commands between the CFE and a selected sequencer, the minimum control latency is five clocks as shown below in Table 2.

Table 4: programmed latency = minimum + 1 (4)

		0	1	2	3	4	5	6	7
5	Command Decoder		D1	D2	D3	D4			
						C1			
	Functional Unit					command issue			
						~ C2 C3 C4			
		↑				↑			
10		Command Received				First control action			

Table 5: programmed latency = minimum + 2 (5)

		0	1	2	3	4	5	6	7	8
15	Command Decoder		D1	D2	D3	D4				
						command issue				
	Functional Unit					~ C1 C2 C3 C4				
		↑				↑				
20		Command Received				First control action				

It is also possible to incorporate the an embodiment of the invention within a control system with only a single functional unit. In that case, command processing is broken down into a front end block which performs command decoding and issues commands to a single back end block that executes control actions. The decomposition of the control system into two parts therefore allows parallelism even with a single functional unit because the back end block can perform the control actions for command N even as the front end decoder processes command N+1. The invention may be applied as in the case of multiple functional units to reduce minimum control latency.

In general, this invention may be used in any application where it is important to reduce the minimum

latency within a control system processing a stream of commands or instructions where the control actions for two or more separate commands may overlap in time and control latency is programmable. These include high speed pipelined interchip communications interfaces, packet based network routing and bridging equipment, specialized data processing and digital signal processing systems, and control and stimulus generation within automated test equipment (ATE).

The improvements attained through the implementation of the present invention include a reduction in the minimum control action latency compared to the conventional scheme with a front end decoder unit issuing commands to multiple parallel functional units and implementation of all control actions within the parallel functional units while achieving the same minimum latency as an aggressive implementation with replicated command decoding logic in each parallel functional unit while avoiding most of its extra complexity and power consumption relative to the conventional scheme.

With respect to the general implementation of the half period adjust scheme, the proposed solution can be used for any application where digital control signals must be generated with a timing resolution too small to be practical or desirable for conventional synchronous control

logic with timing resolution equal to the clock period. This could include high speed interchip communication schemes, output waveform shaping circuits, programmable time base generators, automated test equipment (ATE),
5 direct digital synthesis (DDS) signal generators, and high frequency signal modulation.

The above disclosure is to be taken as illustrative of the invention, not as limiting its scope or spirit. Numerous modifications and variations will become apparent
10 to those skilled in the art after studying the above disclosure. For example, apparatus according to the invention need not issue commands to a sequencer exactly simultaneously with the performance of the first memory control step(s). It is sufficient for the apparatus to
15 issue the command "substantially" simultaneously with the performance of the first memory control step(s), such as within one clock cycle.

As used herein, a given signal or event is "responsive" to, or "depends upon", a predecessor signal
20 or event if the predecessor signal or event influenced the given signal or event. If there is an intervening processing element or time period, the given event or signal can still be "responsive" to, or "dependent upon", the predecessor signal or event. If the intervening
25 processing element combines more than one signal or event,

the signal output of the processing element is considered "responsive" to, or "dependent upon", each of the signal or event inputs. If the given signal or event is the same as the predecessor signal or event, this is merely a degenerate case in which the given signal or event is still considered to be "responsive" to, or "dependent upon", the predecessor signal or event.

Given the above disclosure of general concepts and specific embodiments, the scope of protections sought is to be defined by the claims appended hereto.